

AD-R179 338

COMPILE SMALLTALK-80 TO A RISC(U) CALIFORNIA UNIV  
BERKELEY DEPT OF COMPUTER SCIENCES W R BUSH ET AL.  
06 AUG 87 N00034-K-0251

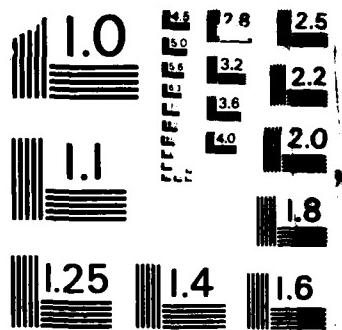
1/1

UNCLASSIFIED

F/G 9/2

NL





MICROCOPY RESOLUTION TEST CHART  
NATIONAL BUREAU OF STANDARDS - 1963 -

AD-A179 330

Productivity Engineering in the UNIX† Environment

DTIC FILE COPY

Compiling Smalltalk-80 to a RISC

Technical Report

S. L. Graham  
Principal Investigator

(415) 642-2059



"The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government."

Contract No. N00039-84-C-0089

August 7, 1984 - August 6, 1987

Arpa Order No. 4871

†UNIX is a trademark of AT&T Bell Laboratories

DISTRIBUTION STATEMENT A  
Approved for public release  
Distribution Unlimited

874  
C17

# **Compiling Smalltalk-80 to a RISC**

*William R. Bush*  
*(principal author)*

Division of Computer Science  
571 Evans Hall  
University of California at Berkeley  
Berkeley, CA 94720  
(415) 642-4694

## *A. Dain Samples*

Division of Computer Science  
University of California at Berkeley

*David Ungar*

Electrical Engineering Department  
Stanford University

## *Paul N. Hilfinger*

Division of Computer Science  
University of California at Berkeley

AD-1179330

## REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION unclassified		1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION/AVAILABILITY OF REPORT unlimited	
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE			
4. PERFORMING ORGANIZATION REPORT NUMBER(S)		5. MONITORING ORGANIZATION REPORT NUMBER(S)	
6a. NAME OF PERFORMING ORGANIZATION The Regents of the University of California	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION SPAWAR	
6c. ADDRESS (City, State, and ZIP Code) Berkeley, California 94720		7b. ADDRESS (City, State, and ZIP Code) Space and Naval Warfare Systems Command Washington, DC 20363-5100	
8a. NAME OF FUNDING/SPONSORING ORGANIZATION DARPA	8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER	
8c. ADDRESS (City, State, and ZIP Code) 1400 Wilson Blvd. Arlington, VA 22209		10. SOURCE OF FUNDING NUMBERS	
		PROGRAM ELEMENT NO.	PROJECT NO.
		TASK NO.	WORK UNIT ACCESSION NO.
11. TITLE (Include Security Classification) Compiling Smalltalk-80 to a RISC			
* 12. PERSONAL AUTHOR(S) W.R. Bush, A.D. Samples, D. Ungar, P. Hilfinger			
13a. TYPE OF REPORT technical	13b. TIME COVERED FROM _____ TO _____	14. DATE OF REPORT (Year, Month, Day) * _____	15. PAGE COUNT * 11
16. SUPPLEMENTARY NOTATION			
17. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP	
19. ABSTRACT (Continue on reverse if necessary and identify by block number)  Enclosed in paper.			
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS		21. ABSTRACT SECURITY CLASSIFICATION unclassified	
22a. NAME OF RESPONSIBLE INDIVIDUAL		22b. TELEPHONE (Include Area Code)	22c. OFFICE SYMBOL

- A

## Compiling Smalltalk-80 to a RISC

### ABSTRACT

The Smalltalk On A RISC project at U.C. Berkeley proved that good performance executing a high-level language on a modified RISC architecture was attainable. The system removed a layer of interpretation, translating bytecodes to a register-based instruction set. This paper describes the bytecode translator and how it was affected by SOAR architectural features. The translator generates code of reasonable density and speed. Because of bytecode semantics, relatively few optimizations are possible, but hardware and software mechanisms at runtime offset these limitations. Register allocation with register windows is the major task of the compiler. Several hardware features could be easily performed by the compiler.



Accession For	
NTIS CRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution / _____	
Availability Codes	
Dist	Avail and/or Special
A-1	

## Compiling Smalltalk-80 to a RISC

### 1. Introduction

The goal of the Smalltalk On A RISC (SOAR) project at U.C. Berkeley was to produce a high-performance execution engine for the Smalltalk-80<sup>1</sup> language [1]. The heart of the effort is a Berkeley RISC processor extended to support Smalltalk [2]. The processor was designed in conjunction with the runtime system [3], which together have yielded substantial performance improvements over conventional Smalltalk-80 implementations. An extensive performance evaluation can be found in Ungar's dissertation [4]; SOAR runs Smalltalk roughly 2.5 times as fast as a Motorola 68010 with a similar cycle time, which works out to be about the same speed as the Xerox Dorado high-performance workstation. One reason for the SOAR system's speed is its compilation of Smalltalk. It has been estimated that compiling Smalltalk to the SOAR instruction set produces a factor of two speedup over conventional interpreted systems. This paper describes the SOAR compiler, presenting the mechanisms of compiling Smalltalk to the special-purpose SOAR RISC instruction set.

### 2. The Nature of the Compiler

The Smalltalk-80 language is defined operationally in terms of a virtual machine that executes stack-based instructions called bytecodes. The Smalltalk-80 programming environment is a binary image that runs on the virtual machine. The problem with the virtual machine is that it is inefficient if naively implemented. It is commonly realized as a bytecode interpreter, which requires special hardware (such as that possessed by the Dorado) to avoid interposing a layer of overhead between the virtual machine and the native machine.

An alternative is to compile bytecodes to native machine instructions, an approach successfully taken by Deutsch and Schiffman [5]. The Deutsch and Schiffman system dynamically translates procedures (called methods in Smalltalk) as needed, keeping a cache of native-code methods and flushing the least-recently-used ones. The

---

<sup>1</sup>Smalltalk-80 is a trademark of Xerox Corporation.

SOAR system takes this approach a step further and compiles all methods from bytecodes into SOAR instructions -- the canonical representation of a method in the SOAR system is as SOAR instructions, not bytecodes. Where caching is done for space efficiency, compiling everything simplifies and speeds up the system for a moderate cost in space.

We considered compiling Smalltalk directly to SOAR, avoiding bytecodes entirely, but we did not take this path for three reasons. First, the virtual machine is the semantic definition of the language. This implies that a correct bytecode compiler is, along with correct system functions, a correct Smalltalk implementation. Second, keeping bytecodes as an intermediate form permits mixed-mode debugging like that found in LISP systems, where functions to be debugged are interpreted, and tested functions that should run fast are compiled. Third, the standard virtual image is a bytecode-based image; to produce a SOAR-based native image those bytecodes must be translated to SOAR instructions. Rather than develop two separate compilers, one for taking bytecodes to SOAR and another for compiling Smalltalk to SOAR, we implemented one bytecode compiler with two implementations: a C version for image conversion, and a Smalltalk version for use with the converted image. Debugging is addressed in more detail in Lee's report [6], and image conversion in a recent paper [3].

The basic task of the compiler is to translate stack-oriented bytecodes into RISC-style loads, stores, and other register-based instructions. It does this by assigning Smalltalk variables and stack locations to registers and memory locations, and then simulating at compile time the bytecode stack operations, converting them to SOAR operations. The simulated stack is used to remember value sources and operations; when a value destination is encountered the code to load, compute, and store the value is generated. If the Smalltalk variables A and B are assigned to registers, for example, a push of A and a pop-and-store into B is translated into a register-to-register move; the simulated stack is used to remember the source A until the destination B is encountered. The stack is simulated at compile time to avoid unnecessary computations at runtime.

In this regard, the Smalltalk bytecodes perform the same function as any stack-based intermediate language such as, for example, UCODE [7]. Bytecodes are unlike UCODE in that, as we will discuss later, they restrict the compiler writer, particularly in implementing optimizations that would result in code motion or that require type information to perform. For example, common subexpression elimination is exceptionally difficult to do in Smalltalk.

Given the restrictive semantics of the Smalltalk bytecodes and the simple architecture of SOAR, the compiler falls or stands on its success in mapping Smalltalk variables to registers and memory.

### 3. SOAR Register Windows

One feature of the Berkeley RISC architectures is a register file of overlapping register windows, each window corresponding to a procedure activation frame. The windows are allocated on procedure call in a stack discipline, using the registers in the window overlap to pass parameters. The advantage of register windows is fast procedure call and return, avoiding the saving and restoring of registers. Tests with benchmarks indicate that SOAR Smalltalk would be 46% slower without them.<sup>2</sup>

It is crucial that the size of these register windows be chosen wisely. If a method requires more storage than can fit in a window, the extra values are spilled to memory, slowing down the procedure call and the method's execution as well as increasing its code size. If windows are too small, an excessive number of methods will spill, degrading the performance of the whole system. On the other hand, if windows are too large, registers will be wasted and fewer windows can be accommodated on the processor chip.

For SOAR, the goal was to have 90% to 95% of all activation frames fit in SOAR windows. Preliminary studies [9] indicated that a window size of 16 registers with complete overlap was best. The SOAR windows are smaller than those of other RISC designs such as RISC II, which has 32, because Smalltalk methods are

---

<sup>2</sup> All results quoted in this paper are found in either Ungar's dissertation [4] or Bush's report [8], unless stated otherwise.

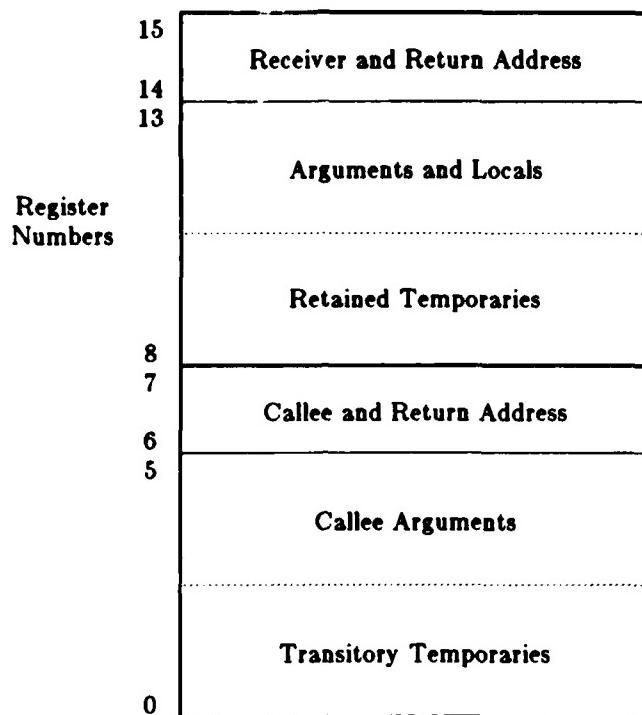
correspondingly smaller than procedures in more traditional imperative-style languages.

An alternative would have been to use variable size windows at some expense in hardware [10]. For SOAR, we concluded that variable size windows would not result in significant enough improvement to offset the additional hardware complexity.

The SOAR hardware divides each window into two identical sets of 8 registers, a high set (15-8) and a low set (7-0). Each set contains 6 general purpose registers and 2 dedicated registers used for return addresses (15 and 7) and return values (14 and 6).

The current method receives its parameters and stores its local variables in the high register set. It sets up parameters for any methods that it calls in the low set. Unlike other RISC designs, the SOAR architecture has no registers dedicated solely for local use -- all are shared between two activation records. This allows the compiler flexibility in register allocation -- registers can be used for arguments or local values depending on what is appropriate. Temporaries that must persist through calls to other methods ('retained' temporaries) are put in free high registers. 'Transitory' temporaries (those whose life spans do not cross procedure calls, such as intermediate results from compiler-generated expressions) can be put in the lows. Figure 1 presents this categorization pictorially.

Allocating variables and temporaries to registers is complicated by the fact that it is possible to write a Smalltalk method whose variables and temporaries will not all fit in a register window. When more registers are required than are available, we resort to spilling. There are two rules used to determine what and when to spill to memory. The first rule of assignment by category specifies that entire categories of variables are spilled -- if not all of the arguments fit in the registers, for example, all are spilled. The second rule of permanent assignment means that a variable cannot be moved once it has been allocated a location -- if, for example, a local variable has been put in a register, it will not later be spilled to make room for a temporary. Neither of these rules result in minimal register usage or minimal memory traffic, but they are reasonable and simple. Since a major goal of the SOAR architectural design was to minimize spills at a reasonable cost, their infrequent occurrence justified an easily



**Figure 1: Register Allocation in a SOAR Register Window**  
 modifiable, simple spilling strategy.

The details of our allocation and assignment strategy are straightforward. The compiler assigns up to four arguments and locals to a method's high registers, and the remaining high registers are used for retained temporaries. If there are more arguments than will fit in four registers, all arguments are put in memory. If the arguments will fit in four registers, but the local variables will not fit with the arguments, all of the locals are stored in memory. Transitory temporaries can be stored in the lows. Retained temporaries are stored in registers if any are available at that point in the computation, and in memory otherwise.

A nice benefit of the above rules is that for most Smalltalk methods, code can be generated in one pass over the bytecodes by making some initial assumptions about the method's register requirements. If the assumptions turn out to be false, a second pass can be made with the new information. Only 9% of all methods in the

Smalltalk-80 system require the second pass.<sup>3</sup>

Spills can be implemented in one of two ways. First, space can be allocated from a common spill pool, or a separate spill object can be allocated for each activation frame that spills. The former has complications for garbage collection and processes, and the latter eats up a register in each spilling frame. Both techniques have been tried in SOAR; neither has shown itself to be clearly superior over the other.

#### 4. How the SOAR Architecture Simplifies the Compiler

Register windows and spills complicate the compiler; SOAR architectural support for Smalltalk has in the main simplified it. An important architectural feature allows the compiler to generate standard arithmetic and comparison instructions in spite of the fact that the operations may be on non-integer objects.

Since Smalltalk is polymorphic and variable types are not known at compile time, it is never safe to generate integer instructions without runtime type checking. The virtual machine requires a dynamic method lookup on each operator, including '+'. Implementing this lookup naively is very expensive. Studies have shown [11] that in fact most arithmetic operations in Smalltalk involve only integers. SOAR takes advantage of this fact by assuming all simple arithmetic operations (plus, minus, comparisons, etc.) will be performed only on integers. The compiler thus treats all such operations as if they were on integers, and generates integer code. If, say, an add instruction is initiated on two objects, and one or both of them turn out to be non-integers, the hardware will trap and transfer to a handler that will look up the correct method for the intended operation. When the execution of the looked-up method is complete, the trap handler returns to the instruction immediately following the one that caused the trap.

It would not be difficult for the compiler to generate code to test explicitly for integer tags, but it would slow down system performance by an estimated 20% and

---

<sup>3</sup>The Smalltalk and the C versions of the compiler differ in this regard: the C version effectively makes two passes over each method. However, speed was not as critical for the C version as for the Smalltalk version.

would increase image size by 15%.

The tag mechanism is also used to provide hardware assist for garbage collection, thus further simplifying the task of code generation. SOAR uses a generation scavenging scheme, dividing memory into regions of old and new objects, and uses generation tags on pointers to detect when pointers to new objects are stored into old ones. The tag check is performed in the hardware, and a trap handler records the necessary information if a trap is taken. Benchmark results indicate that tagged stores are so rare that the compiler could in fact generate explicit tests and only slow the system about 1% and expand the image 2%.

The compiler also takes advantage of hardware that maps registers to memory addresses, and allows what we call a pointer to register. Since a Smalltalk program can access any object in memory, and activation records are just another kind of object, it is necessary to handle references to them. Remember that activation records exist as on-chip register windows. The pointer to register feature permits the compiler to ignore complications that would otherwise be caused by overt references to activation records. Once again, however, it turns out that checking for pointers to activation records is not as severe a problem as it was initially feared to be. The instances requiring such a check do not occur often, and L'ngar concluded that the feature could be removed with only a 3% performance penalty.

## 5. How Bytecode Compilation Simplifies the Compiler

Both language and pragmatic considerations limited the optimizations we could perform.

First, Smalltalk is a polymorphic language. That is, the same 'A+B' expression within a Smalltalk method could on one instantiation add two integers and on the next concatenate two strings. In the latter case it is not true that 'A+B' equals 'B+A'. Furthermore, which method is invoked is, by definition, a function of the leftmost operand ('A' in our example). Together these facts preclude almost all expression evaluation optimization.

Second, the fact that activation frames are full-fledged Smalltalk objects requires a relatively straightforward mapping between bytecode frames and compiled ones, eliminating optimizations such as method integration.

Third, our pragmatic goal was to bring up a working Smalltalk-80 image for the SOAR processor. This made us fundamentally conservative in our approach to the language. Correctness, as defined by the bytecodes and the image, was more important than efficiency. Because of this we were loath to change the image (including the standard Smalltalk-80 compiler) or take advantage of ambiguities in the language specification. We were not in the normal position of a compiler writer.

There is one optimization which has been employed successfully by Deutsch and Schiffman and the SOAR system. While it is true that the language permits polymorphic expressions, the fact of the matter is that 98% of the time the method invoked does not change. Thus, for example, if 'A+B' added complex numbers the last time it was evaluated, it is highly likely that it will be used to add complex numbers the next time it is evaluated. Using this observation we cache the last method called. When it is called again, we merely check that the type of the leftmost operand this call matches the type of the leftmost operand from the last call. If they do not match the standard method lookup mechanism is re-invoked. Ungar estimated that removing this caching strategy would slow down our system by 33%.

## 6. Results

Three types of performance measurements were made with the SOAR compiler: compilation speed, code expansion, and register window utilization.

Since the SOAR compiler is an extra stage added to the bytecode compiler, it can only slow compilation down. However, on a Dorado (equivalent in performance to SOAR) the compiler is reasonably fast both objectively and subjectively. Results from compiling the entire Smalltalk-80 image of 4770 methods indicate that it adds a mean time of 50 milliseconds to a method's total compilation time. This lengthens compilation by 70%. Subjectively, the compiler does not intrude on system use since it is usually invoked interactively on one method at a time. A mean total compilation time of about 120 milliseconds per method does not noticeably affect response time.

One of the virtues of bytecodes is that they are a compact representation. A major concern at the start of the project was an anticipated explosion in method size that could result from moving to four-byte word-sized instructions. Preliminary estimates indicated potential expansion of up to 1000% [12]. Fortunately, observed expansion with the actual implementation is considerably lower. For the entire Smalltalk-80 image the mean bytecode method size is 32.5 bytes; that expands to a mean length of 40.6 SOAR words,<sup>4</sup> an expansion factor of 499%. We note here that the largest method compiles into 637 bytes of bytecodes; it takes the Smalltalk-80 compiler and our compiler approximately 3 seconds to compile this method from Smalltalk into 684 SOAR words. These expansion results are very similar to the expansion factor of 503% reported by Deutsch and Schiffman [5] for their comparable translator with in-line cache.

To put these averages in perspective, total method storage required for the bytecode image is 155025 bytes, and is 774648 bytes for the SOAR version. For current workstations with memories in the four to eight megabyte range, this increase of roughly 600k bytes is not significant.

A major goal of the SOAR design was to have over 90% of the methods fit all their storage into registers. Static results from the compiler support the current window size. Only 9% of all methods spill. Examining total register requirements also supports the 16 register window. These numbers are expressed in terms of the numbers of high registers used, because all method-specific allocation is done in the high registers.<sup>5</sup>

number of high registers	methods using no more than that number
2	29%
4	65%
8	92%
16	99%
32	100%

Dynamic results from benchmarks confirm the static results, and more emphatically

---

<sup>4</sup>The method lengths include literals and in-line cache instructions.

<sup>5</sup>These numbers include the 2 special receiver and return address registers. Note that the percentage of methods that spill is higher than the percentage of methods that require more than 8 registers because the registers are divided up into two regions (arguments and locals, and temporaries) that have different spill disciplines.

endorse the chosen window size. Dynamically, less than 3% of methods spill. The above figures show that if we increased the register window set size to 32 from 8, all methods in the Smalltalk-80 system would fit in register windows, and spilling could effectively be eliminated. However, since less than 3% of all calls require spilling, even if we could (somehow) keep the same number of register windows on the chip, increasing the window size from 16 to 64 registers would not be cost effective.

Static results also indicate that the current spill rules are reasonable.

category	mean size	maximum size
arguments	0.83	13
local variables	0.68	19
retained temporaries	0.82	10
total registers	2.32	28
spill area	0.45	22

Argument, local variable, and temporary demands are small on the average, but in the worst case all exceed the window size.

## 7. Conclusions

The Smalltalk-80 language and its bytecode representation restrict the conventional optimizations available to the SOAR compiler. Nonetheless, the compiler generates efficient code, primarily due to register windows, integer tags and traps, and in-line method caches. Experience with the compiler has verified the architectural design decision to use a 16-register fully overlapped window. Several features supported by the SOAR hardware could easily be performed by the compiler at marginal increased time and space costs.

## 8. Acknowledgments

This project was sponsored by Defense Advance Research Projects Agency (DoD) ARPA Order No. 3803, monitored by Naval Electronic System Command under Contract No. N00034-K-0251. It was also sponsored by Defense Advance Research Projects Agency (DoD) Arpa Order No. 4871, monitored by Naval Electronic Systems Command under contract N00039-84-C-0089. Dain Samples is supported in part by an AT&T Bell Laboratories Scholarship. Xerox PARC provided many machine cycles in the development of the compiler.

### References

- [1] *Smalltalk-80 -- The Language and its Implementation*; Adele Goldberg and David Robson; Addison-Wesley, 1983.
- [2] 'Architecture of SOAR: Smalltalk on a RISC'; David Ungar, Ricki Blau, Peter Foley, A. Dain Samples, David Patterson; 11th Annual International Symposium on Computer Architecture, Ann Arbor, Michigan, June 1984.
- [3] 'SOAR: Smalltalk Without Bytecodes'; A. Dain Samples, David Ungar, Paul Hilfinger; *CACM Object-Oriented Programming Systems, Languages and Applications Proceedings*, November 1986, pp 107-118.
- [4] 'The Design and Evaluation of a High Performance Smalltalk System'; David Michael Ungar; *U.C. Berkeley Computer Science Division Report UCB/CSD 86/287*, March 1986.
- [5] 'Efficient Implementation of the Smalltalk-80 System'; L. Peter Deutsch, Allan M. Schiffman; *11th ACM Symposium on Principles of Programming Languages*, January 1984, pp 297-302.
- [6] 'The Design of a Debugger for SOAR'; Peter K. Lee; Master's Report, U.C. Berkeley, 1984.
- [7] 'Machine Independent Pascal Code Optimization'; D.R. Perkins and R.L. Sites; *CCC SIGPLAN 14,8*, August 1979, pp 201-207.
- [8] 'Smalltalk-80 to SOAR Code'; William R. Bush; *U.C. Berkeley Computer Science Division Report UCB/CSD 86/297*, June 1986.
- [9] 'Register Windows for SOAR'; John Blakken; *Smalltalk On A RISC -- Proceedings of CS929R*, April 1983, pp 126-140.
- [10] 'Reduced Instruction Set Computer Architectures for VLSI'; Manolis G.H. Katevenis; *U.C. Berkeley Computer Science Division Report UCB/CSD 88/141*, October 1983.
- [11] *Smalltalk-80 Bits of History, Words of Advice*; Glenn Krasner, Editor; Addison-Wesley, 1983.
- [12] 'Implementing a Smalltalk Compiler'; Wayne Citrin, Carl Ponder; *Smalltalk On A RISC -- Proceedings of CS929R*, April 1983, pp 167-185.

E N D

5 — 8 7

D T / C